

# Graphics Chip Makes Low-Cost, High Resolution Color Displays Possible

by Mark Olson and Brad May

The making of displays that are both high-resolution and low-cost is the key to producing equipment for both the automated office and the engineering workstation. Through the introduction of 16-bit  $\mu$ Ps such as Intel's iAPX 8088, 80186 and 80286, the processing power has been made available to perform very sophisticated functions for the user while making the human interface very simple.

That processing power can be unnecessarily drained, however, if the  $\mu$ P is burdened with the entire task of graphics display. Such a burden can fill up a significant part of the processor's I/O bandwidth, slow down the refresh rate of the display, and decrease the computational power of the CPU.

## Intelligent peripheral ICs offload processing tasks from the CPU.

The logical way to avoid such limitations is to dedicate a specialized processor to the handling of display function. It should be capable of accepting high-level commands to minimize the burden on the CPU, as well as optimizing the execution of such commands through raster operations imple-

mented in hardware at the device level.

Such a chip is Intel's 82720 Graphics Display Controller (GDC). It has features that give systems a fast drawing speed while reducing graphics display costs by 60% or more. It achieves these results by taking over the drawing and refresh functions from the CPU, by allowing the use of dynamic RAM's instead of static RAM's, and by reducing the overall parts count needed to create a complete graphics system.

The implementation of the drawing task is a major feature of the GDC. Other graphics chips perform only the display refresh function, leaving the more complicated drawing function entirely to the CPU. Since the CPU is doing every pixel of the drawing function on these systems, they also require fas-

ter bit map RAM than with the GDC. The GDC, on the other hand, is capable of handling the drawing function itself, drawing such objects as characters, slanted characters, points, lines, arcs, rectangles, and slanted rectangles based only upon lengths, slopes, and arc centers supplied by the CPU. The GDC's processing, moreover, takes place concurrently with the processing of the CPU.

### 2048 x 2048 Resolution

With its 4 Megapixel addressability, one GDC can handle a monochrome display with resolution as high as  $2048 \times 2048$ , and multiple GDC's can be linked to provide even higher resolution, such as color displays at  $2048 \times 2048$ . The chances are, however, that the GDC's full power will not be used in most applications. The typical

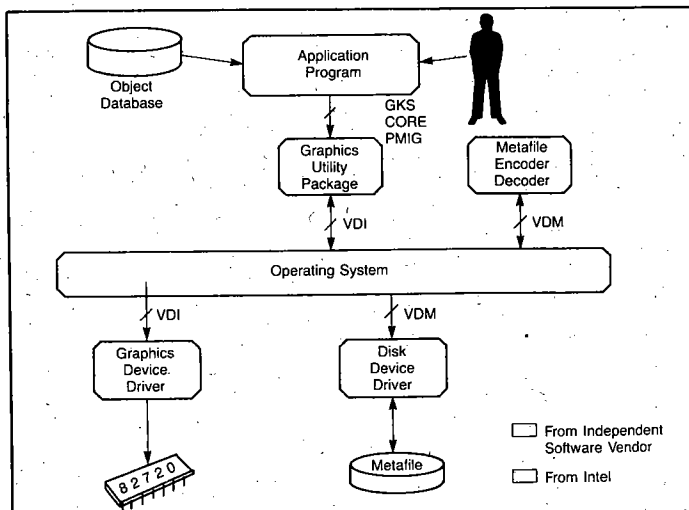


Figure 1: General graphics commands are translated into the VDI interface level and then into driver device commands.

Mark Olson and Brad May are Product Marketing Engineers for Peripheral Components Operation, Intel Corp., Santa Clara, CA 95051.

Reprinted from DIGITAL DESIGN © April 1983, Morgan-Grampian Publishing Company, Boston, MA 02215

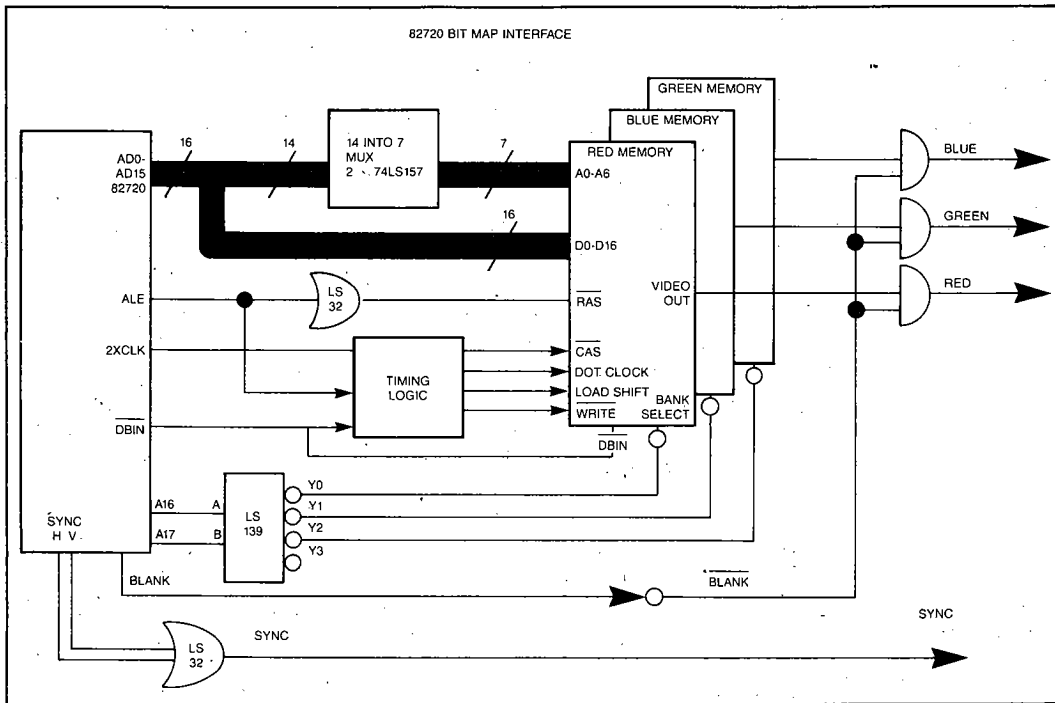


Figure 2: The memory is broken up into three planes, with each plane feeding one of the primary color guns of the CRT.

product considered high resolution for office automation applications is a  $512 \times 512$  pixel monochrome or color display.

These latter restrictions are not imposed by the GDC, but rather have more to do with the cost of display monitors, the amount of RAM memory needed to support such displays, and the adequacy of such displays for most applications. It is possible to build "super graphics" boards with a GDC, such as the 1K by 1K pixel by 8 color plane graphics display designed by Phoenix Computer Graphics (Lafayette, LA). Such a display is capable of rendering 256 different colors on a high resolution screen.

Even higher performance can be achieved through the use of multiple GDC's to support multiple display windows, increased drawing speed, or increased bits per pixel. For multiple display windows, each GDC can be used to control one window of the display. For in-

creased drawing speed, multiple GDC's can be operated in parallel. For increased bits/pixel, each GDC can contribute a portion of the number of bits necessary for a pixel.

Although the GDC is intended primarily for raster-scan graphics, it can also be used as a character display controller. It is capable of supporting up to four screens of data containing 25 rows by 80 columns, or one screen containing up to 100 rows by 256 characters.

### Office Automation Display

High performance applications can stretch the usage of the GDC from low-end to high-end engineering displays, but research has shown that for office automation products, a  $512 \times 512$  pixel display is quite acceptable, and that color is often a requirement. These requirements mesh with a major factor in display—the cost of the CRT. In

OEM quantities, for example, one could expect to find a  $512 \times 512$  monochrome display for under \$100, a  $256 \times 256$  color display (TV quality) for about \$150, a  $512 \times 512$  color CRT in the \$300 range, and a  $1K \times 1K$  color display in the \$800–\$1000 category.

To give an example of the type of display that can be built for new office products using the GDC, consider a  $512 \times 512$  pixel by 3 color plane combination CPU and graphics display on a single 12" by 12" board. Such a display is capable of generating 8 colors.

The list of parts (Table 2) comes to about \$175 for 85 IC's taking up 104 square inches of board space. Even that parts count could be reduced by replacing the 48 16K DRAMs with 12 64K DRAMs—if a  $4K \times 16$  bit DRAM were available. A very important note about the parts list is that the design is implemented with inexpensive 2118 dynamic RAMs. The design does

not require the faster, more expensive, and less dense static RAMs.

The parts count is low enough so that the processor and graphics controller can be placed together in a single 12" by 12" board. This is important because small overall size and footprint are selling points for desktop workstations. System speed is also enhanced when the graphics controller and CPU are on the same board, because their communication need not take up bus, inter-board bandwidth or experience any additional delays.

### Pipelining Transformations

More important than putting the graphics display on the same board

as the CPU is the level of communication between the CPU and graphics controller. If the burden of transformation processing is left entirely to the CPU while the graphics chip is used only as a CRT controller, then the CPU must communicate one bit per pixel to update a display. With the GDC, the CPU input takes higher level forms such as the slope and length of a line, the length and center point of an arc, or the key coordinates of a rectangle. Since the average line on a screen is about 25 pixels, that means that 25 times fewer CPU bus cycles are required to draw a graphical object with the GDC. These CPU cycles (an average of 50  $\mu$ s each to calculate the graphical object and communicate it to

the GDC) are the determining factor in drawing rate.

Viewed from a larger perspective, there are four tasks that must be performed by a CPU-graphics chip combination:

(1.) The CPU must calculate the higher-level graphics operations. This is done by the CPU and it involves the processing of macro-operations such as the CORE, GKS, PMIG or other graphics protocols. These general graphics commands are translated into an intermediate level, the VDI interface level (**Figure 1**) and then into device driver commands by software in the CPU.

(2.) Then, these lower-level graphical objects such as the key parameters for lines, arcs, characters, and rectangles, must be trans-

## VLSI Takes Aim At Text Processing

The concept of co-processing is not a new one. Intended as a way of offloading computationally intensive tasks from a host CPU, it has been around at Intel since the introduction of the 8087 numerics processor and the 8089 I/O machine. A more recently developed product, the 82720 Graphics Display Controller is designed to bolster system performance by offloading graphics control chores from the CPU. The chip accepts high level commands from the CPU and, using its own drawing processor, accesses the required positions in the bit-map and handles the processing and display control functions.

Building on the success of these parts come two new co-processors designed to partition system intelligence even further. The 82586 is a communications co-processor designed to bridge the characteristics of CPU and network data rates. Its FIFO buffer and DMA facilities make it possible for a CPU to operate at the full Ethernet 10 Mbits/s transfer rate even in the face of continuous bursts of network data traffic.

Intel's most recent introduction is the 82730 text co-processor. Printers and other hard copy peripherals have supported additional text processing features such as proportional spacing and simultaneous superscript and subscript for some time. Implementing these features on the display screen has traditionally been a costly procedure. Thus, it is typically not done and screen displays often are not identical to their hard-copy printouts. Aimed to solve this designers headache, the 82730 has its own DMA capability and communicates asynchronously with the CPU via shared memory messages. It supports the generation of high quality text displays through features like proportional spacing, simultaneous superscript/subscript, dynamically reloadable fonts and user programmable field and character attributes. In addition, when coupled with the 82720 Graphics Display Controller (**Figure 1**) the 82730 provides flexible mixing of text and graphics simultaneously on the same display.

—Wilson

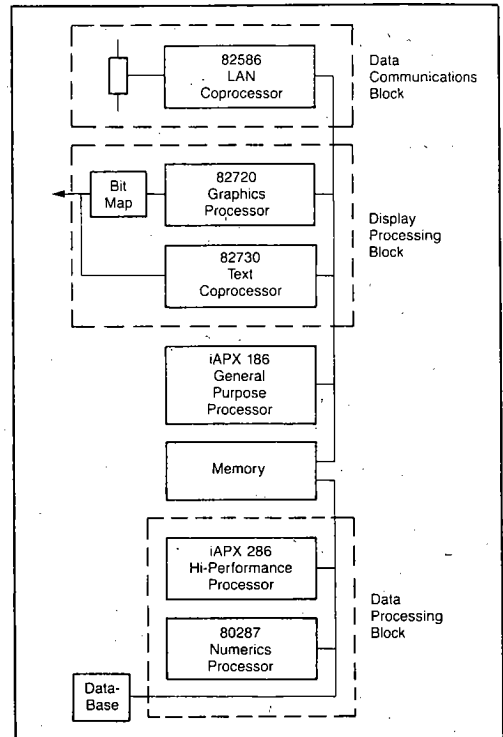


Figure 1: Offloading system tasks is simplified by new VLSI devices.

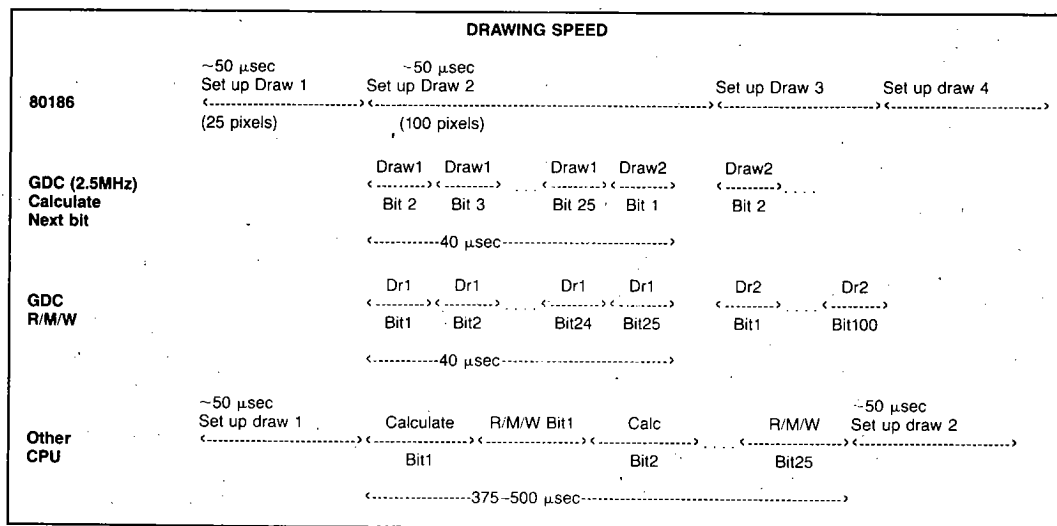


Table 1: The 80186 and the GDC work together to accomplish the drawing function.

formed into changes in the actual bits. This function is performed in hardware in the GDC concurrently with any level one processing done by the CPU. Other graphics controllers leave this task to the CPU to execute in software. The contrast is that, in such systems, the CPU must resolve the graphical object down to every point on a line, while with the GDC it need only designate the endpoints.

(3.) With the actual bits for the bit map calculated, they must be placed in the bit map memory. This involves a read-modify-write operation that requires three CPU cycles using other methods. With the GDC these operations are not the responsibility of the CPU. The GDC pipelines its execution so that it is calculating the next bit to change while it is executing the read-modify-write cycles.

(4.) Finally, the bit map memory must be dumped into the CRT. This is the refresh function performed by other graphics chips as well as the GDC.

The summation is that other systems require the CPU to process steps one to three serially, leaving only step four for the graphics controller. Systems with the GDC require the CPU to process only step one, with the GDC concurrently

processing steps two through four. The GDC has another advantage in that during the transformation process in step three, the GDC executes the algorithms in hardware while a CPU must execute the algorithms in software. The algorithms are exactly the same in both cases. They are the Bresenham algorithms from IBM, in which the next pixel to be drawn becomes a binary decision between two pixels.

The execution of these algorithms is a crucial drawing time factor, because they are invoked many times for each updated screen. Consider that, in the inner loop of Bresenham's "line drawing algorithm," there are two or three additions, two comparisons or tests, and the masking of the proper value into the word for each pixel. The algorithms for drawing circles or filling areas are even more complex. In the inner loop of a fill algorithm, the old word must be read from the bit map, then tested to see if all, some, or none of the pixels are within the area to be filled. Next, it tests whether some or all of the pixels must be modified. Finally, the word must be returned to the bit map.

These algorithms are heavily used and the speed with which they can be executed has a direct effect

upon the overall system efficiency. If they must be executed by a  $\mu$ P, the instruction fetching process slows down the calculations to a drawing rate of 15–20  $\mu$ s per pixel. With a hardware implementation of these algorithms in the GDC, the calculations can be speeded up to achieve a drawing rate of 1600 ns (2.5 MHz version) or 800 ns (5 MHz version) per pixel.

## Methods Of Refresh

In the fourth step, the dumping of bit map memory into the CRT, there are some differences between graphics controller chips. Motorola's MC6845 CRT controller, for example, uses a split-cycle refresh method in which each refresh cycle is alternated with a drawing cycle in which the  $\mu P$  updates the bit map. This gives the MC6845 a 50% drawing bandwidth.

With the GDC there are two drawing modes. The first is a "draw anytime" mode which replaces CRT refresh cycles with drawing cycles. This is the fastest mode, but it does result in on-screen disruptions. The second mode, which does not disrupt the on-screen display, draws only during the vertical and horizontal retracing periods. This gives the GDC about a 25%

1 80186	1 74LS04	1 20 MHz Clock
1 82720	1 74LS73	2 27128
2 74LS157	9 74LS244	2 2186
1 74LS139	8 74LS166	1 8274
1 74LS161	3 74LS32	1 8042
1 74LS11	2 8286	3 Connectors
1 74LS00	1 8 MHz Crystal	1 12 x 12 2 Layer PC

**SUMMARY:**

4 VLSI Controllers	80186 Processor
	82720 Graphics
	8274 Serial Link
	8042 Keyboard
4 VLSI Memory	27128 EPROM
	2186 IRAM
48 16K DRAMs	2118 DRAM
29 MSI/SSI	— Buffers/Glue

**TOTAL:** 85 IC'S → 104 Sq. Inches  
Parts Cost → About \$175

Table 2: Parts list for 512 x 512 x 4 Color Display.

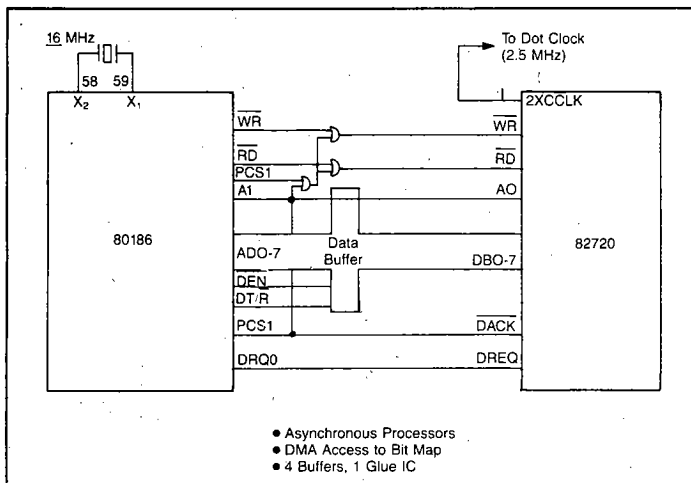


Figure 3: The two chip selects are OR'd together to qualify the R/W signals.

drawing bandwidth. At first glance that gives the GDC a disadvantage in drawing rate, but the fact is, with its pipelining and hardware execution of transformations, the GDC makes much more efficient use of its bandwidth. The critical timing factor is the amount of CPU participation in the drawing process, not the refresh bandwidth of the graphics controller. Another tradeoff is that, with its split-cycle architecture, the MC6845 requires RAM memory that is twice as fast as that required by the GDC in the

same application.

## Inexpensive RAM Is Fast Enough

Applying this perspective, one can begin to build the display with parts listed in Table 2. First one notes that a square display, as indicated by the 512 x 512 pixel initial specification, is not pleasing to the eye. It is much more appealing to have an aspect ratio of about 4:3, in which the number of pixels horizontally is 4/3 the number vertically. If the resolution is such that the total num-

ber of pixels is not a power of two, it will be necessary to round up to the next power of two and waste the extra bits.

The pixel arrangement which best meets this requirement is one with a 432 x 576 pixel format. It also meets the requirement that the number of pixels horizontally be an even number of 16-bit words. With three color bits per pixel (red, blue, and green), the total display memory is then about 500 x 500 x 3, or 750k bits.

It makes the most sense to break the memory up into three planes, with each plane feeding one of the primary color guns of the CRT (Figure 2). This leads to a memory arrangement of 16K x 16 x 3, using 16K dynamic RAMs with a 1K x 16 architecture. When drawing graphics figures, the memory can be treated as one large plane, split into the three primary colors. Drawing in low-order memory could represent red, middle-order could be used for green, and high-order for blue.

One advantage of this 3D memory is that drawing with a primary color requires setting only one bit per pixel. Drawing with a secondary color such as cyan, yellow, or magenta would take two GDC cycles, and creating white from all three colors would take three GDC cycles. If this were an issue, additional hardware could be used to draw more than one plane at a time. As the results will show, however, the drawing speed requirements can be exceeded without any added hardware.

## Calculate The Drawing Rate

To see if the proposed design is practical, one should first calculate the drawing rate to see what the user interface will be like. Then one should check the refresh rate to make sure the design is uninterrupted and without flicker.

The proof of the assumption that CPU participation is the dominating factor lies in the 50 μs average time that it takes the CPU to calculate a graphical object and communicate its key parameters to the GDC. Assume that the graphical object is an average line containing

The GDC's normal clock rate is 2.5 MHz, giving it a 400 ns period (the maximum clock rate is 5 MHz, with a 200 ns period.) It takes four GDC cycles to execute a read-modify-write on a bit (because two read cycles are required), so that the GDC's normal drawing rate is one pixel per 1600 ns. To draw the 25 pixels involved in the average line, then, would take  $25 \times 1600$  ns, or 40  $\mu$ s. Since this operation is done concurrently with CPU processing, the GDC will be waiting for the next graphical object by the time the CPU is ready.

If the screen were filled with nothing but 25-pixel vectors, then the drawing rate would be determined by the 50  $\mu$ s average CPU calculation and transfer cycle, averaging about 2  $\mu$ s per pixel. If all the vectors were white (worst case), then it would take 1.5 secs of drawing time to update the white screen. Since, in the undisturbed-screen mode, drawing is only done

In reality, however, the screen will not be filled with vectors. It will have an average of 500 vectors, and the color distribution could be presumed to be evenly distributed as one-third primary colors, one-third secondary colors, and one-third white. The 500 vectors will require the drawing of 12.5K pixels in monochrome, or 25K pixels with distributed colors. At a drawing rate of 2  $\mu$ s per pixel, this takes 50 ms to draw. Drawing only during blanking, the screen would be updated in 200 ms.

Under these conditions, it would not help to use the maximum clock rate GDC (5 MHz), but if in some applications the average vector length is 100 pixels, then the CPU calculation-and-bus cycle (50  $\mu$ s) would remain the same and the GDC's drawing cycle (1600 ns  $\times$  100 = 160  $\mu$ s) would become a limiting factor. Using the 5 MHz GDC would cut that drawing time

down to 800 ns/pixel, or 80  $\mu$ s/vector. The 500 vector average screen would then contain 100K pixels with distributed colors and could be drawn in 80 ms. Multiplying by four because the drawing is done during blanking (25% of the time), that is 320 ms. That is a screen update in less than one-third second for a "busy" screen.

## Calculate The Refresh Rate

These calculations are of little importance if the display flickers due to lack of refresh. This exercise is actually a demonstration of how the basic GDC clock rate was derived. Assume a non-interlaced display that must be refreshed 60 times per second. That gives a screen refresh rate of 16.67 ms, but on a typical CRT some 4.27 ms of that is blanked, leaving 12.4 ms of active display time. The dot sweep period is the 12.4 ms divided by the number of pixels ( $432 \times 576 = 248.8\text{K}$ ), or 49.8 ns. The inverse gives a 20.07 MHz dot clock.

Since the GDC dumps 16 bits from the bit map memory into the

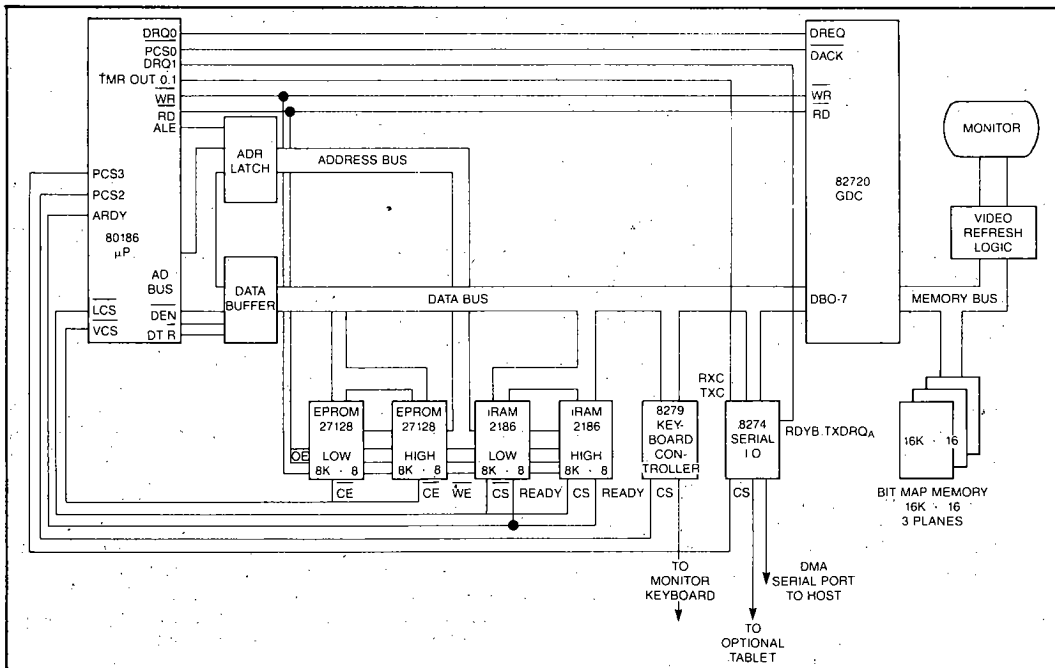


Figure 4: Completed graphics system uses the 80186 and 82720 GDC.

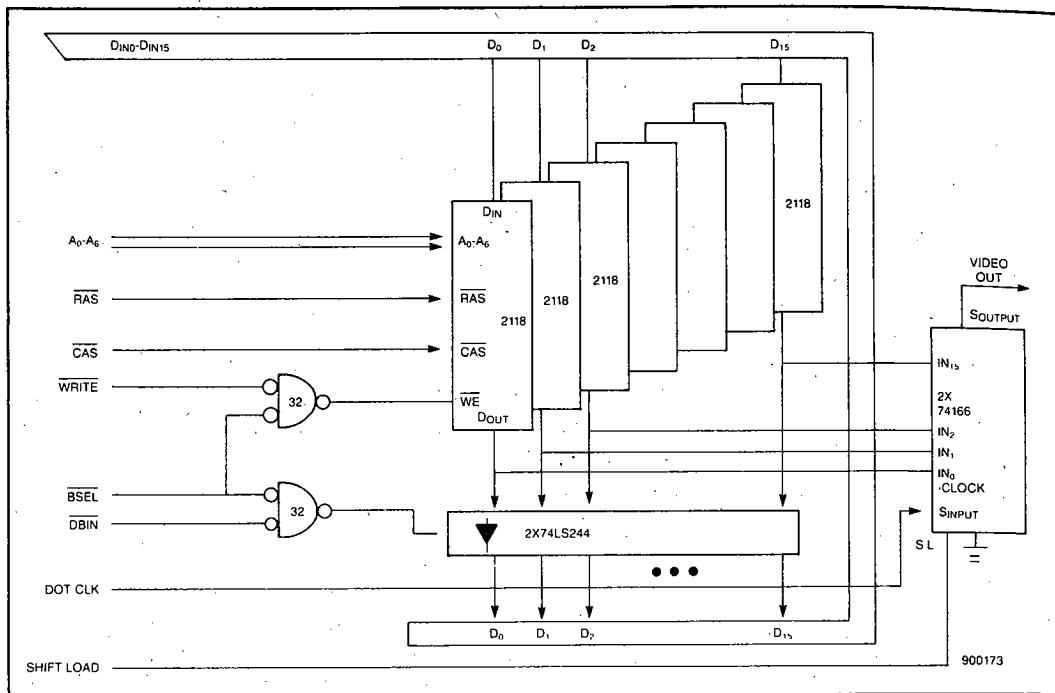


Figure 5: Since the 186 is a fully byte addressable machine, it is possible to write bytes as well as words into the RAMs.

16-bit shift register during each read, and since the shift register then feeds these bits out serially to the CRT, it makes sense that the GDC's read period should be 16 times the dot sweep period. That gives a GDC read period of about 800 ns. With each GDC read taking two cycles, the basic GDC clock period is then 400 ns, or 2.5 MHz. This gives a rock-solid display, and one would only want to go to the 5 MHz GDC to improve drawing rate.

For those who want to examine the blanking intervals to see if the CRT is indeed "typical," the blanking can be further broken down. The vertical blanking interval is 1.25 ms, leaving 15.42 ms to scan the 432 lines on the active portion of the display. Dividing 15.42 ms by 432 lines gives a 35.7  $\mu$ s period per line, or a horizontal sweep rate of 28 KHz. Time is also needed for horizontal retrace, in this case, 7  $\mu$ s of horizontal blanking per line. This leaves 28.7  $\mu$ s to scan the 576

pixels on each line, resulting in the dot sweep period of 49.8 ns. Using a 20 MHz CRT helps keep the costs down, but the GDC can use CRT displays as fast as 80 MHz when higher resolution is required.

### Mixed Mode

While it is possible to generate 8 $\times$ 8 characters and slanted characters in the graphics mode, the GDC also offers a mixed mode memory organization to display both characters and graphics drawn from separate windows in the display memory. The advantage of this mode is that it allows characters to be manipulated as 8-bit entities instead of the 64 bits that each would require in graphics mode. Of necessity, the graphics window display memory is reduced in this mode (64K 16-bit words instead of 256K), but even the reduced maximum graphics memory is still a megapixel and quite sufficient for both office automation and engineering display purposes.

In the character window, the GDC operates as it does in the pure character mode, with the exception that the line counter must be implemented externally. In addition to the two windows used for graphics and characters in the mixed mode, two other windows can be supported. These can be designated as either character or graphics windows by a selection on the A17 line.

### Panning, Zooming, Light Pen

As special features, the GDC allows both panning and zooming in either graphics, character, or mixed modes. The zoom is accomplished by effectively increasing the size of the dots on the screen. Vertically, this is done by repeating the same display line. The number of repeat times is determined by the display zoom parameter. Horizontally, zoom is accomplished by extending each display word cycle and displaying fewer words per line, according to the zoom factor.